On tensor rank of conditional probability tables in Bayesian networks

Jiří Vomlel[1] and **Petr Tichavský** [2]

arXiv:1409.6287v1 [cs.AI] 22 Sep 2014

Abstract. A difficult task in modeling with Bayesian networks is the elicitation of numerical parameters of Bayesian networks. A large number of parameters is needed to specify a conditional probability table (CPT) that has a larger parent set. In this paper we show that, most CPTs from real applications of Bayesian networks can actually be very well approximated by tables that require substantially less parameters. This observation has practical consequence not only for model elicitation but also for efficient probabilistic reasoning with these networks.

1 INTRODUCTION

On the most cumbersome tasks in modeling with Bayesian networks is the elicitation of conditional probabilities. These probabilities can be estimated from collected data or elicited with the help of domain experts. A general conditional probability table (CPT), $P(X_i|pa(X_i))$, where X_i is a discrete random variable and $pa(X_i)$ are its parent variables, requires $n_i = (|X_i| - 1) \cdot \prod_{X_j \in pa(X_i)} |X_j|$ parameters. The cardinality $|X|$ of a variable X is the number of values the variable X can take. Since n_i is exponential with respect to number of parents the task of probability elicitation can be quite demanding for tables with a large parent set.

To ease the task of probability elicitation special types of CPTs were proposed. They include, so called, canonical models [2] whose typical examples are noisy-or, noisy-max and noisy-threshold, etc. For these special types of CPTs the number of parameters is often linear or quadratic with respect to the number of parents. Furthermore, it is possible to perform more efficient computations with them [5, 3, 7, 8].

On the other hand the vast majority of Bayesian networks used in real applications that are available in Bayesian network repositories uses only general CPTs. The main motivation of this paper is to show that actually many of these general tables have a hidden simple structure that requires a lower number of parameters than general CPTs. We will apply the CP-tensor decomposition [4, 1] for CPTs [7, 8] of models from a BN repository. We will see that many of the CPTS of these Bayesian networks can be approximated pretty well by tables with a low rank.

This observation has two practical consequences. First, more attention should be paid to discovering the internal structure of CPTs that are either estimated from data or elicited with the help of domain experts. Second, the internal structure of CPTs can be exploited for more efficient probabilistic inference with these networks.

[1] Institute of Information Theory and Automation of the AS CR, Prague, Czech Republic, email: vomlel@utia.cas.cz

[2] Institute of Information Theory and Automation of the AS CR, Prague, Czech Republic, email: tichavsk@utia.cas.cz

The rest of this paper is organized as follows. In Section 2 we briefly introduce the CP tensor decomposition. The results of experiments with the state-of-the-art numerical algorithms for the CP-tensor decomposition are reported in Section 3. We conclude the paper by a summary in Section 4.

2 CP TENSOR DECOMPOSITION

Each CPT can be viewed as a tensor. A tensor is simply a mapping $\mathcal{A} : \mathbb{I} \rightarrow \mathbb{R}$, where $\mathbb{I} = I_1 \times \ldots \times I_k$, k is a natural number called the order of tensor $\mathcal{A}$, and $I_j, j = 1, \ldots, k$ are index sets. Typically, I_j are sets of integers of cardinality n_j. Then we can say that tensor $\mathcal{A}$ has dimensions $n_1, \ldots, n_k$. Tensor $\mathcal{A}$ has rank one if it can be written as an outer product of vectors:

$$\mathcal{A} \quad = \quad \boldsymbol{a}_1 \otimes \ldots \otimes \boldsymbol{a}_k$$

with the outer product $\otimes$ being defined for all $(i_1, \ldots, i_k) \in I_1 \times \ldots \times I_k$ as $\mathcal{A}_{i_1,\ldots,i_k} = a_{1,i_1} \cdot \ldots \cdot a_{k,i_k}$, where $\boldsymbol{a}_j = (a_{j,i})_{i \in I_j}$, $j = 1, \ldots, k$ are real valued vectors.

CP-tensor decomposition [4, 1] is decomposition of a tensor into a sum of r tensors of rank one. The minimum number r such that the decomposition is possible is called the rank of tensor $\mathcal{A}$.

In [7] the CP tensor decomposition (called tensor rank-one decomposition there) was applied to CPTs of canonical models. See [7] or [8] for a more detailed explanation, which is out of scope of this paper. Instead, we will just illustrate savings one may get with the CP tensor decomposition of CPTs. Assume a CPT defined for a binary variable with $k - 1$ binary parents. To fully specify such a table we need 2^{k-1} parameters. If this table can be approximated by a table with rank r we need $k(r-1)+r$ parameters instead. If $r = 2$ then we need $k + 2$ parameters. We can immediately see that instead of an exponential number of parameters we need just a linear number of them.

3 NUMERICAL EXPERIMENTS

For the experiments we used 15 Bayesian networks from a Bayesian network repository[3], namely:

alarm.net, hailfinder.net, barley2.net, pathfinder.net, munin1.net, munin2.net, munin3.net, munin4.net, mildew.net, hepar2.net, andes.net, win95pts.net, water.net, link.net, and insurance.net.

In the experiments we considered all CPTs for variables having at least three parents. It was 492 CPTs in total. For the CP tensor de-

[3] We used networks from the repository at `http://www.bnlearn.com/bnrepository/`, where networks from other repositories after a quality-check and fixing are stored.

composition we used the Fast Damped Gauss-Newton (Levenberg-Marquard) algorithm [6] implemented in the Matlab package TENSORBOX [4]. In the first experiment we used this algorithm with 100 different random starting points and one precomputed (using nvec method). In the second experiment we used 10+1 starting points.

First, we present the results for an example of a $3 \times 3 \times 4 \times 3$ CPT from the *hailfinder.net* network. In this CPT the child variable *Boundaries* (3 states) has three parents *OutflowFrMt* (3 states), *MorningBound* (3 states), and *WndHodograph* (4 states). In Figure 1 we compare the maximum distance

$$\max_{(i_1,\ldots,i_k)} |\hat{\mathcal{A}}_{i_1,\ldots,i_k} - \mathcal{A}_{i_1,\ldots,i_k}|$$

of the approximation $\hat{\mathcal{A}}$ to the original CPT $\mathcal{A}$ as a function of the rank of the approximation. We can see that the distance rapidly decreases with rank (full line), which is not the case of a table of the same dimensions with random values (dashed line).

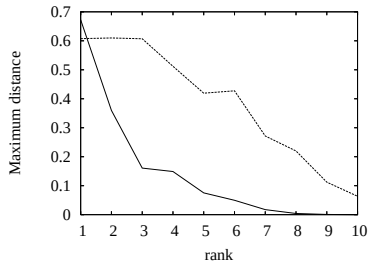

Figure 1. Maximum distance of the approximation to the original CPT as a function of its rank.

Next, we summarize results for all 492 tables from 15 different Bayesian networks. In Figure 2 we present the plot describing dependence of percentage of CPTs that can be approximated by tables of a given rank with maximum error (taken over all possible configurations of parents) smaller than 0.001. See the full line. From this figure we can see that about 40% of the CPTs can be pretty well approximated by tables of rank 2, about 75% by tables of rank 6, and only about 15% require rank more than 10 to get a very good fit.

One may conjecture that any CPT can be approximated well by a table with a similarly low rank as CPTs of networks from repositories. To show that this is not the case we repeated the computations for CPTs of the same dimensions but with their entries replaced by random numbers. See dashed line in Figure 2. We can see that it is not possible to approximate randomly created CPTs with similarly low rank as for CPTs of Bayesian networks from the repository.

4 CONCLUSIONS

We presented numerical experiments on 15 Bayesian network models from real applications. The results suggest that most of the conditional probability tables can be very well approximated by tables that have lower rank than one would expect from a general table of the same dimensions. Our results are in line with results presented in [9] where the authors experimentally verified that noisy-max provides a surprisingly good fit for as many as 50% of CPTs in two of the three Bayesian networks they tested.

[4] Tensorbox is freely available at `http://www.bsp.brain.riken.jp/~phan/tensorbox.php`.

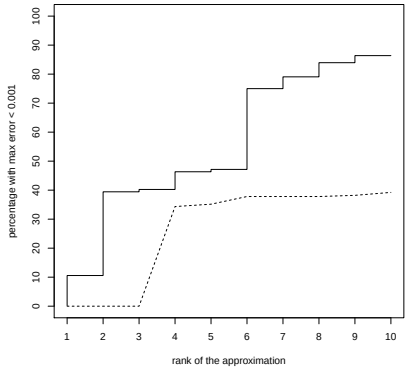

Figure 2. Percentage of CPTs that can be approximated with a maximum error smaller than 0.001 as a function of the rank of the approximation.

The low rank approximation should be exploited not only in the model elicitation by using a compact parametrization of the internal structure of a CPT but also during the probabilistic inference. We conjecture that some low rank approximations of CPTs may actually correspond better to what a domain expert intended to model in the constructed CPT. One reason might be that when filling numerical values into a CPT with more than two parents it is quite demanding for the model creator to provide exact numerical values. Another reason might be that numerical values in the CPTs are often provided with a lower numerical precision, typically with two decimal points.

ACKNOWLEDGEMENTS

This work was supported by the Czech Science Foundation through projects No. 13–20012S and No. 14–13713S.

REFERENCES

[1] J. D. Carroll and J. J. Chang, 'Analysis of individual differences in multidimensional scaling via an n-way generalization of Eckart-Young decomposition', *Psychometrika*, **35**, 283–319, (1970).

[2] F. J. Díez and M. J. Druzdzel, 'Canonical probabilistic models for knowledge engineering', Technical Report CISIAD-06-01, UNED, Madrid, Spain, (2006).

[3] F. J. Díez and S. F. Galán, 'An efficient factorization for the noisy MAX', *International Journal of Intelligent Systems*, **18**, 165–177, (2003).

[4] R. A. Harshman, 'Foundations of the PARAFAC procedure: Models and conditions for an "explanatory" multi-mode factor analysis', *UCLA Working Papers in Phonetics*, **16**, 1–84, (1970).

[5] K. G. Olesen, U. Kjærulff, F. Jensen, F. V. Jensen, B. Falck, S. Andreassen, and S. K. Andersen, 'A MUNIN network for the median nerve — a case study on loops', *Applied Artificial Intelligence*, **3**, 384–403, (1989). Special issue: Towards Causal AI Models in Practice.

[6] A.-H. Phan, P. Tichavský, and A. Cichocki, 'Low complexity damped Gauss-Newton algorithms for CANDECOMP/PARAFAC', *SIAM Journal on Matrix Analysis and Applications*, **34**, 126–147, (2013).

[7] P. Savicky and J. Vomlel, 'Exploiting tensor rank-one decomposition in probabilistic inference', *Kybernetika*, **43**(5), 747–764, (2007).

[8] J. Vomlel and P. Tichavský, 'Probabilistic inference with noisy-threshold models based on a CP tensor decomposition', *International Journal of Approximate Reasoning*, (2013).

[9] A. Zagorecki and M. J. Druzdzel, 'Knowledge engineering for Bayesian networks: How common are noisy-MAX distributions in practice?', *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, **43**, 186–195, (2013).